

Claims

What is claimed is:

1. A method of creating an index for a database table of records, the method occurring in a computer environment having a plurality of processing units wherein each processing unit has access to the table, the method comprising:

determining partition delimiters, each partition delimiter separating the table into non-overlapping partitions of records, each partition dedicated to one processing unit for index creation;

accessing the table records in parallel, wherein each processing unit accesses each of the records;

filtering the accessed records in parallel, wherein each processing unit determines which records to keep;

independently creating a plurality of sub-indexes, wherein at least two sub-indexes are created by different processing units; and

merging the sub-indexes together to create a final index related to the table.

2. A method as defined in claim 1 wherein the act of creating the sub-indexes further comprises sorting the records and generating a data structure based on the sorted records.

3. A method as defined in claim 2 wherein the data structure is a B-Tree data structure.

4. A method as defined in claim 2 wherein the data structure has multiple levels.

5. A method as defined in claim 2 wherein the data structure is a clustered index.

6. A method as defined in claim 1 further comprising gathering sub-index statistical information and stitching sub-index statistical information.

7. A method as defined in claim 1 wherein the method is initiated by an index creation manager module.

8. A method as defined in claim 1 wherein the method is initiated by a query manager in response to a supplied query.

9. A method as defined in claim 1 wherein the method is initiated automatically in response to a modification to the table.

10. A method as defined in claim 1 wherein the act of determining partition delimiters comprises:

sampling the table records to determine an approximate distribution of the values in the key field;

creating a histogram based on the sampled information; and

evaluating the histogram to determine the partition delimiters.

11. A method as defined in claim 10 further comprising:

determining a processor goal value based on the number of processors in the computer system;

determining a least common multiple value based on the processor goal value;

determining whether the histogram information may be substantially evenly split into the least common multiple value number of partitions;

if so, creating the partition delimiters based on the least common multiple value;

and

if not, adjusting the processor goal to determine a new least common multiple value to determine partition delimiters.

12. A computer program product readable by a computer and encoding instructions for executing the method recited in claim 1.

13. A computer program product readable by a computer and encoding instructions for executing the method recited in claim 11.

14. A system for database table index creation for a database table, the database table stored in memory and comprising a plurality of records, the system comprising:

a plurality of processing units that respectively accesses the database table in parallel, the respective processing units accesses each of the records and filters the accessed records to determine which records to keep and wherein the respective processing units creates a sub-index of database table records; and
a merge tool that merges the plurality of sub-indices into a final database table index.

15. A system as defined in claim 14 wherein each processing unit further comprises:

a scanning module that scans the database table;
a filter module that filters the accessed records and selectively predetermined records; and
a sorting module that sorts records kept by the filter module into a sub-index.

16. A system as defined in claim 15 wherein the scanning module, filter module and sorting module, for each processing unit, operate concurrently.

17. A system as defined in claim 15 further comprising a sampling module for sampling the database table and a partition module for dividing the records into substantially equal quantities related to the number of processing units.

18. A method of creating an index for a database table of records, the method occurring in a computer environment having a plurality of processing units wherein more than one processing unit has access to the table, the method comprising:

determining partition delimiters, each partition delimiter separating the table into non-overlapping partitions of records, wherein at least one partition is dedicated to a first processing unit for index creation and at least one other partition is dedicated a second processing unit for index creation;

the first processing unit accessing a table record and determining whether the table record is associated with the at least one partition dedicated to the first processing unit; and

the first processing unit only processing the accessed table record when the accessed table record is associated with the at least one partition dedicated to the first processing unit.

19. A method as defined in claim 18 further comprises:

upon determining that the accessed table record is not associated with the at least one partition dedicated to the first processing unit, passing the accessed record to the second processing unit for index creation.

20. A method of creating an index for a database table of records, the method occurring in a computer environment having a plurality of processing units wherein more than one processing unit has access to the table, the method comprising:

determining partition delimiters, each partition delimiter separating the table into non-overlapping partitions of records, each partition dedicated to one processing unit for index creation;

independently creating a plurality of sub-indexes, wherein at least two sub-indexes are created by different processing units;

allocating blocks of a disk to store each sub-index, wherein parts of each sub-index are stored on consecutive blocks on the disk; and

merging the sub-indexes together to create a final index related to the

table.

21. A method as defined in claim 20 wherein the act of allocating portions of the disk allocates a predetermined number of blocks, the predetermined number of blocks is determined during the determination of the partition delimiters.

22. A method as defined in claim 20 wherein the allocation of portions of the disk comprises:

maintaining a cache of allocated pages and allocating pages for each partition in the cache for each processing unit; and
retrieving a pre-determined number of database pages upon request, and wherein the number of pages to allocate upon each request is determined by the size of the cache.

23. A method as defined in claim 22 wherein the cache has a size depending on the size of the index being built and the number of currently available free pages in the system.

24. In a computer system having a plurality of processors, an index creation system for creating an index of information for a table of data records, the system comprising:

a sampling module that samples the table of data records to determine sub-index delimiters;
two or more index creation modules, each index creation module associated with a processor, each index creation module creates a sub-index; and
a merge module that merges the sub-indexes into a final index,
wherein each index creation module comprises:
an access module that accesses data records from the table of data records;
a filter module that filters data records according the sub-index

delimiters to keep only relevant data records; and
a sorting module that sorts the relevant data records into a sub-index.

25. A system as defined claim 24 further comprising a memory allocation module that allocates parts of memory for storing the sub-indexes, and wherein the memory allocation module allocates a predetermined number of parts, the predetermined number of parts is determined during the determination of the delimiters.

26. A system as defined in claim 24 further comprising a cache memory module that manages a cache of allocated pages and allocates pages for storing each sub-index in the cache and wherein the number of pages allocated to the cached is determined upon determining the delimiters.

27. An index creation system for creating an index of information for a table of data records, the system comprising:

- means for sampling the table of data to determine sub-index delimiters;
- means for accessing data records from the table in parallel;
- means for filtering accessed data records to keep only relevant records;
- means for creating two or more sub-indexes of relevant records; and
- means for merging the sub-indexes together.

28. An index creation system as defined in claim 27 further comprising:
means for allocating memory for storing parts of each sub-index in contiguous memory blocks.